arXiv:1403.7683v1 [math.ST] 30 Mar 2014

Approximate Matrix Multiplication with Application to Linear Embeddings

Anastasios Kyrillidis
Computer and Communication Sciences, EPFL
anastasios.kyrillidis@epfl.ch

Michail Vlachos
IBM Research Lab, Zurich
mvl@zurich.ibm.com

Anastasios Zouzias*
IBM Research Lab, Zurich
azo@zurich.ibm.com

April 1, 2014

Abstract

In this paper, we study the problem of approximately computing the product of two real matrices. In particular, we analyze a dimensionality-reduction-based approximation algorithm due to Sarlos [1], introducing the notion of *nuclear rank* as the ratio of the nuclear norm over the spectral norm. The presented bound has improved dependence with respect to the approximation error (as compared to previous approaches), whereas the subspace – on which we project the input matrices – has dimensions proportional to the maximum of their nuclear rank and it is independent of the input dimensions.

In addition, we provide an application of this result to linear low-dimensional embeddings. Namely, we show that any Euclidean point-set with bounded nuclear rank is amenable to projection onto number of dimensions that is independent of the input dimensionality, while achieving additive error guarantees.

1 Introduction

Living in the era of *Big Data*, the excess of available information constitutes its manipulation and interpretation a strenuous task: In contrast to conventional wisdom where more data is a source of "simplicity" in statistical terms [2], large datasets embody preprocessing tasks of high time- and space-complexity, jeopardizing any hope for data analysis within reasonable time and with low computational cost. Due to such difficulties, one might be interested in *screening* the data, even if *no prior information is available*. I.e., identifying a coreset such that most of the latent structure/information that we want to infer is maintaned within an error that we can control. Such approaches have been witnessed in a broad class of data analysis problems such as data matrix sparsification for accelerated spectral calculation [3], feature selection (a.k.a. column subset selection) for better interpretation of the results [4] and low-complexity calculations [5], feature extraction via dimensionality reduction techniques [6, 7], etc.

In this work, we analyze a particular approximation algorithm of [1] for the task of matrix multiplication with respect to their spectral norm [8]. Within this context, we focus on approximately computing the product of two matrices *when their intrinsic dimensionality is low, as expressed by the nuclear rank.* Here, we define the nuclear rank as the ratio of the nuclear norm of a matrix over its spectral norm. We provide an elementary proof based on the randomized algorithm described in [1, 8]. As a result, we further strengthen the performance of the proposed scheme in situations where the allowed approximation error is small: Using a weaker notion of intrinsic dimensionality (i.e., nuclear rank instead of stable rank; see the definitions later in text), the dependence on the approximation error ε is improved to $\mathcal{O}(1/\varepsilon^2)$ instead of $\mathcal{O}(1/\varepsilon^4)$, indicating that the proposed scheme scales better when ε decreases. Table 1 places our result into context with prior works.

Table 1: Summary of results on ε-approximate matrix multiplication $\mathsf{A}\cdot\mathsf{B}$ with respect to spectral norm – $\mathbf{nr}(\cdot)$ and $\mathbf{sr}(\cdot)$ denote the nuclear and stable rank of a matrix, respectively.

Metric	**# of dimensions**	**Reference**
Rank	$\mathcal{O}((r(\mathsf{A})+r(\mathsf{B}))/\varepsilon^2)$	[9]
Stable rank	$\mathcal{O}((\mathbf{sr}(\mathsf{A})+\mathbf{sr}(\mathsf{B}))/\varepsilon^4)$	[8]
Nuclear norm	$\mathcal{O}((\mathbf{nr}(\mathsf{A})+\mathbf{nr}(\mathsf{B}))/\varepsilon^2)$	Theorem 1

As an application, we use this result to design dimensionality reducing linear embeddings that operate on a given dataset and preserve Euclidean distances between data points. In general, it is well-known that simply rotating, scaling and translating in a random way is adequate for this task: According to JL lemma [10], such linear mapping approximately preserves (i.e., within a ε-radius of points) the distances between all the pairs of points. While there are many random constructions that

*The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement n^o 259569.

achieve this property with high probability, such schemes are oblivious to the geometry of the data set at hand, leaving space for further improvements.

Contributions: The main contributions of this manuscript are the following:

(i) We provide a novel analysis on the approximate matrix multiplication problem based on the notion of nuclear rank (See Theorem 1).

(ii) We demonstrate an application of the approximate matrix multiplication bound to dimensionality reduction with relative guarantees.

2 Related Work

There is rich literature on the topic of linear dimensionality reduction and the techniques utilized for this purpose, mainly due to the wide range of applications it covers. Here, we highlight a few approaches that provide bounds depending on the input data: These bounds are usually stronger than the classic JL lemma, provided that the input points have low "geometric complexity", i.e., points that lie on a low dimensional subspace [1], lie on a manifold [11, 12], etc. Here, we focus only on results that provide theoretical guarantees.

For the special case where the input points lie on the unit sphere, there is a close connection between the Talagrand functional, denoted as γ_2, and the required number of dimensions for point distance preservation [13],[14]. These results can be viewed as stronger bounds compared to the JL lemma. Baraniuk et al. provide bounds on the number of dimensions required by a random linear embedding to preserve Euclidean distances for the case of smooth manifolds [11]; for improvements and a nice exposition on the topic, see [12]. Sarlos provided improved bounds for algorithms on large matrices with applications on matrix multiplications, linear regression and low rank matrix approximation [1]. On a similar flavor, the authors in [15] provide bounds for relative error approximations of points that lie on a surface. Several approaches exist for non-linear dimensionality reduction as well [16, 17]. From a different perspective, Indyk and Naor consider the "doubling dimension" and other related measures of dimensionality for preserving nearest neighbor queries [18].

Recent developments in [19] describe *deterministic constructions* of linear embeddings in polynomial time via SemiDefinite Programming (SDP) relaxations, under the assumption that the data is known *apriori* and fixed. For the latest developments on this topic, we refer the reader to [20].

3 Preliminaries

A scalar is denoted by an italic letter, e.g. a. A column vector is denoted by a bold lowercase letter, e.g. $\mathbf{a} \in \mathbb{R}^n$ whose i-th entry is a_i. A matrix is denoted by a mathtype uppercase letter, e.g., $\mathsf{A} \in \mathbb{R}^{n\times d}$ with (i,j)-th entry $\mathsf{A}_{i,j}$. To denote the i-th row and j-th column of A, we use the $\mathbf{a}_i$ and $\mathbf{a}^j$, respectively. Use $\|\mathsf{A}\|_2 = \max_{\mathbf{x}:\|\mathbf{x}\|_2=1} \|\mathsf{A}\mathbf{x}\|_2$ to denote the spectral norm of A and $\|\mathsf{A}\|_{\mathrm{F}} = \sqrt{\sum_i \sum_j \mathsf{A}_{i,j}^2}$ to represent its Frobenius norm.

In our analysis, we utilize the notion of *nuclear norm*. Given a matrix $\mathsf{A} \in \mathbb{R}^{n\times d}$ with rank $r(\mathsf{A})$, its nuclear norm $\|\mathsf{A}\|_\star$ is given by $\|\mathsf{A}\|_\star = \sum_{i=1}^{r(\mathsf{A})} \sigma_i$ where σ_i is the i-th largest singular value; for multiple matrices, we also use ξ_i to denote a singular value. Using the nuclear norm, we define the *nuclear rank* of A as $\mathbf{nr}(\mathsf{A}) := \|\mathsf{A}\|_\star / \|\mathsf{A}\|_2$. Moreover, we use the notion of *stable rank*, which is defined as $\mathbf{sr}(\mathsf{A}) = \|\mathsf{A}\|_{\mathrm{F}}^2 / \|\mathsf{A}\|_2^2$.

Throughout the paper, we will denote by G a $t \times d$ random matrix whose entries are independent Gaussian random variables with variance $1/t$, i.e., $\mathsf{G}_{ij} \sim \mathcal{N}(0, 1/t)$.

4 Main result

We consider the following problem:

APPROXIMATE MATRIX MULTIPLICATION PROBLEM: Let $\mathsf{X} \in \mathbb{R}^{n\times d}$ and $\mathsf{Y} \in \mathbb{R}^{d\times m}$ be two arbitrary matrices, $\varepsilon > 0$ is an approximation parameter and $0 < \delta < 1$ the failure probability. We desire to construct *sketched* matrices $\widehat{\mathsf{X}} \in \mathbb{R}^{n\times t}$ and $\widehat{\mathsf{Y}} \in \mathbb{R}^{t\times m}$ where $t \ll d$ such that:

$$\left\|\widehat{\mathsf{X}}\widehat{\mathsf{Y}} - \mathsf{X}\mathsf{Y}\right\|_2 \leq \varepsilon \|\mathsf{X}\|_2 \cdot \|\mathsf{Y}\|_2 ,$$

holds with probability at least $1 - \delta$.

We now state the main theorem of the paper:

Theorem 1. *Fix $0 < \varepsilon, \delta < 1$ and assume arbitrary matrices $\mathsf{X} \in \mathbb{R}^{n\times d}$ and $\mathsf{Y} \in \mathbb{R}^{d\times m}$. Set $\widehat{\mathsf{X}} = \mathsf{X}\mathsf{G}^\top$ and $\widehat{\mathsf{Y}} = \mathsf{G}\mathsf{Y}$. If $t = \Omega\left(\frac{\boldsymbol{nr}(\mathsf{X})+\boldsymbol{nr}(\mathsf{Y})+\log(\log(1/\varepsilon))+\log(1/\delta)}{\varepsilon^2}\right)$, then the following holds:*

$$\mathbb{P}\left(\left\|\widehat{\mathsf{X}}\widehat{\mathsf{Y}} - \mathsf{X}\mathsf{Y}\right\|_2 \leq \varepsilon \|\mathsf{X}\|_2 \cdot \|\mathsf{Y}\|_2\right) \geq 1 - \delta, \tag{1}$$

We devote the rest of this section to prove Theorem 1. We recall a well-established result from the literature for Gaussian matrices; observe the lack of any upper bound on the error parameter.

Lemma 1. *Fix $\zeta > 0$ and let $\mathsf{X} \in \mathbb{R}^{n\times d}$ and $\mathsf{Y} \in \mathbb{R}^{d\times m}$. Set $\widehat{\mathsf{X}} = \mathsf{X}\mathsf{G}^\top$ and $\widehat{\mathsf{Y}} = \mathsf{G}\mathsf{Y}$. Then,*

$$\mathbb{P}\left(\left\|\widehat{\mathsf{X}}\widehat{\mathsf{Y}} - \mathsf{X}\mathsf{Y}\right\|_2 > \zeta \left\|\mathsf{X}\right\|_2 \left\|\mathsf{Y}\right\|_2\right) \leq c_2^{r(\mathsf{X})+r(\mathsf{Y})} e^{-c_1 t\zeta\cdot\min(\zeta,1)}$$

where c_1 is the Hanson-Wright constant [21] and $c_2 = 18$.

Proof. The proof is a corollary of the Hanson-Wright inequality [21, Theorem 1.1], combined with a dense net argument on the unit sphere defined by the union of the column and row span of Y and X, respectively. □

By homogeneity, we observe that, to prove $\left\|\mathsf{X}\mathsf{G}^\top\mathsf{G}\mathsf{Y} - \mathsf{X}\mathsf{Y}\right\|_2 \leq \varepsilon \left\|\mathsf{X}\right\|_2 \left\|\mathsf{Y}\right\|_2$ is satisfied with some probability, it suffices to prove the same argument for $\left\|\frac{\mathsf{X}}{\|\mathsf{X}\|_2}\mathsf{G}^\top\mathsf{G}\frac{\mathsf{Y}}{\|\mathsf{Y}\|_2} - \frac{\mathsf{X}}{\|\mathsf{X}\|_2}\cdot\frac{\mathsf{Y}}{\|\mathsf{Y}\|_2}\right\|_2 \leq \varepsilon$. Thus, without loss of generality, we can assume that $\|\mathsf{X}\|_2 = \|\mathsf{Y}\|_2 = 1$.

Let $\mathsf{X} = \sum_{i=1}^{r(\mathsf{X})} \sigma_i \mathbf{u}_i \mathbf{v}_i^\top$ be the singular value decomposition (SVD) of X, where $\sigma_1, \dots, \sigma_{r(\mathsf{X})}$ denote the singular values of X and, $\mathbf{u}_1, \dots, \mathbf{u}_{r(\mathsf{X})} \in \mathbb{R}^n$ and $\mathbf{v}_1, \dots, \mathbf{v}_{r(\mathsf{X})} \in \mathbb{R}^d$ denote the left and right singular vectors, respectively. Similarly, we define the SVD of Y as $\mathsf{Y} = \sum_{j=1}^{r(\mathsf{Y})} \xi_j \mathbf{p}_j \mathbf{q}_j^\top$.

Define $\theta := \left\lfloor c_3 \cdot \frac{\mathbf{nr}(\mathsf{X})+\mathbf{nr}(\mathsf{Y})}{\varepsilon^2} \right\rfloor$ where $c_3 > 1$ is a constant and set $t := \theta + 8\ln(8/\delta)/\varepsilon^2 + \ln\left(\lceil\ln(e/\varepsilon)\rceil\right)/(c_1\cdot\varepsilon^2)$. Given θ, one can decompose X as $\mathsf{X} = \mathsf{X}_\theta + \mathsf{X}_\theta^c$ where X_θ represents the *best* rank-θ approximation of X and $\mathsf{X}_\theta^c := \mathsf{X} - \mathsf{X}_\theta$. Similarly, we can decompose $\mathsf{Y} = \mathsf{Y}_\theta + \mathsf{Y}_\theta^c$.

Now, by the triangle inequality:

$$\begin{aligned}
\|\mathsf{X}\mathsf{G}^\top\mathsf{G}\mathsf{Y} - \mathsf{X}\mathsf{Y}\|_2 \leq{}& \left\|\mathsf{X}_\theta\mathsf{G}^\top\mathsf{G}\mathsf{Y}_\theta - \mathsf{X}_\theta\mathsf{Y}_\theta\right\|_2 && (2)\\
&+ \left\|\mathsf{X}_\theta^c\mathsf{G}^\top\mathsf{G}\mathsf{Y}_\theta\right\|_2 + \left\|\mathsf{X}_\theta\mathsf{G}^\top\mathsf{G}\mathsf{Y}_\theta^c\right\|_2 + \left\|\mathsf{X}_\theta^c\mathsf{G}^\top\mathsf{G}\mathsf{Y}_\theta^c\right\|_2 && (3)\\
&+ \|\mathsf{X}_\theta^c\mathsf{Y}_\theta\|_2 + \|\mathsf{X}_\theta\mathsf{Y}_\theta^c\|_2 + \|\mathsf{X}_\theta^c\mathsf{Y}_\theta^c\|_2\,. && (4)
\end{aligned}$$

Bounding the term appearing on the right hand side of (2) is the most challenging task and it is the main technical contribution of this paper. Before we start, we need to define some notation. Let l, s be non-negative integers and define the sets $\mathcal{I} = \{\mathcal{I}_1, \mathcal{I}_2, \dots\}$ and $\mathcal{J} = \{\mathcal{J}_1, \mathcal{J}_2, \dots\}$ where:

$$\mathcal{I}_l := \left\{i \;:\; e^{-l} \leq \sigma_i < e^{-l+1}\right\}, \tag{5}$$

and

$$\mathcal{J}_s := \left\{j \;:\; e^{-s} \leq \xi_j < e^{-s+1}\right\}. \tag{6}$$

We remind that, since $\|\mathsf{X}\|_2 = \|\mathsf{Y}\|_2 = 1$, $\sigma_i \leq 1,\ \xi_j \leq 1$, for any i, j.

Given a set $\mathcal{I}_l$, we define $\mathsf{X}_\theta^l \in \mathbb{R}^{n\times d}$ as $\mathsf{X}_\theta^l := \sum_{i\in\mathcal{I}_l} \sigma_i \mathbf{u}_i \mathbf{v}_i^\top$; similarly, we have $\mathsf{Y}_\theta^s := \sum_{j\in\mathcal{J}_s} \xi_j \mathbf{p}_j \mathbf{q}_j^\top$ for the case of Y. We highlight the following key observations:

(i) $\mathsf{X}_\theta = \sum_{l=1}^{L} \mathsf{X}_\theta^l$ and $\mathsf{Y}_\theta = \sum_{s=1}^{S} \mathsf{Y}_\theta^s$, for some bounded positive integers L, S.

(ii) $r(\mathsf{X}_\theta^l) = \mathrm{card}(\mathcal{I}_l)$ and $r(\mathsf{Y}_\theta^s) = \mathrm{card}(\mathcal{J}_s)$.

(iii) $\left\|\mathsf{X}_\theta^l\right\|_2 < e^{-l+1}$ and $\|\mathsf{Y}_\theta^s\|_2 < e^{-s+1}$ by definition of sets $\mathcal{I}_l$ and $\mathcal{J}_s$ respectively in (5)-(6).

In our analysis, we select L and S such that $\|\mathsf{X}_\theta^c\|_2 \leq \varepsilon$ and $\|\mathsf{Y}_\theta^c\|_2 \leq \varepsilon$. By definition of $\mathcal{I}_l,\ \mathcal{J}_s$, one can deduce that $\|\mathsf{X}_\theta^c\|_2 \leq e^{L+1}$ and $\|\mathsf{Y}_\theta^c\|_2 \leq e^{S+1}$, for fixed L and S. To this end, we require $L = \lceil\ln\left(\frac{e}{\varepsilon}\right)\rceil$ and $S = \lceil\ln\left(\frac{e}{\varepsilon}\right)\rceil$ for both $\|\mathsf{X}_\theta^c\|_2 \leq \varepsilon$ and $\|\mathsf{Y}_\theta^c\|_2 \leq \varepsilon$ to be satisfied.

Based on the definitions above and using triangle inequality on $\left\|\mathsf{X}_\theta\left(\mathsf{G}^\top\mathsf{G} - \mathsf{I}\right)\mathsf{Y}_\theta\right\|_2$, we obtain:

$$\begin{aligned}
\left\|\mathsf{X}_\theta\left(\mathsf{G}^\top\mathsf{G} - \mathsf{I}\right)\mathsf{Y}_\theta\right\|_2 &\overset{(i)}{=} \left\|\left(\sum_{l=1}^{L}\mathsf{X}_\theta^l\right)\left(\mathsf{G}^\top\mathsf{G} - \mathsf{I}\right)\mathsf{Y}_\theta\right\|_2\\
&\overset{(iv)}{\leq} \sum_{l=1}^{L}\left\|\mathsf{X}_\theta^l\left(\mathsf{G}^\top\mathsf{G} - \mathsf{I}\right)\mathsf{Y}_\theta\right\|_2\\
&\overset{(i)}{=} \sum_{l=1}^{L}\left\|\mathsf{X}_\theta^l\left(\mathsf{G}^\top\mathsf{G} - \mathsf{I}\right)\left(\sum_{s=1}^{S}\mathsf{Y}_\theta^s\right)\right\|_2\\
&\overset{(iv)}{\leq} \sum_{s=1}^{S}\left(\sum_{l=1}^{L}\left\|\mathsf{X}_\theta^l\left(\mathsf{G}^\top\mathsf{G} - \mathsf{I}\right)\mathsf{Y}_\theta^s\right\|_2\right). \qquad (7)
\end{aligned}$$

$$
\begin{aligned}
\|\mathsf{X}_\theta\left(\mathsf{G}^\top\mathsf{G}-\mathsf{I}\right)\mathsf{Y}_\theta\|_2 &\leq \sum_{s=1}^{S}\sum_{l=1}^{L}\left(\left(\varepsilon + c_2\frac{r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)}{c_1 t}\right)\left\|\mathsf{X}_\theta^l\right\|_2\|\mathsf{Y}_\theta^s\|_2\right)\\
&= \varepsilon\sum_{s=1}^{S}\sum_{l=1}^{L}\left\|\mathsf{X}_\theta^l\right\|_2\|\mathsf{Y}_\theta^s\|_2 + \frac{c_2}{c_1 t}\sum_{s=1}^{S}\sum_{l=1}^{L}r(\mathsf{X}_\theta^l)\left\|\mathsf{X}_\theta^l\right\|_2\|\mathsf{Y}_\theta^s\|_2 + \frac{c_2}{c_1 t}\sum_{s=1}^{S}\sum_{l=1}^{L}r(\mathsf{Y}_\theta^s)\left\|\mathsf{X}_\theta^l\right\|_2\|\mathsf{Y}_\theta^s\|_2\\
&\overset{(iii)}{<} \varepsilon\left(\sum_{s=1}^{S}e^{-s+1}\right)\left(\sum_{l=1}^{L}e^{-l+1}\right) + \frac{c_2}{c_1 t}\left(\sum_{s=1}^{S}e^{-s+1}\right)\left(\sum_{l=1}^{L}r(\mathsf{X}_\theta^l)\left\|\mathsf{X}_\theta^l\right\|_2\right) + \frac{c_2}{c_1 t}\left(\sum_{l=1}^{L}e^{-l+1}\right)\left(\sum_{s=1}^{S}r(\mathsf{Y}_\theta^s)\|\mathsf{Y}_\theta^s\|_2\right)\\
&\leq \varepsilon\left(\sum_{s=1}^{\infty}e^{-s+1}\right)\left(\sum_{l=1}^{\infty}e^{-l+1}\right) + \frac{c_2}{c_1 t}\left(\sum_{s=1}^{\infty}e^{-s+1}\right)\left(\sum_{l=1}^{L}r(\mathsf{X}_\theta^l)\left\|\mathsf{X}_\theta^l\right\|_2\right) + \frac{c_2}{c_1 t}\left(\sum_{l=1}^{\infty}e^{-l+1}\right)\left(\sum_{s=1}^{S}r(\mathsf{Y}_\theta^s)\|\mathsf{Y}_\theta^s\|_2\right)\\
&= \frac{\varepsilon\cdot e^2}{(e-1)^2} + \frac{c_2}{c_1 t}\cdot\frac{e}{e-1}\left[\left(\sum_{l=1}^{L}r(\mathsf{X}_\theta^l)\left\|\mathsf{X}_\theta^l\right\|_2\right) + \left(\sum_{s=1}^{S}r(\mathsf{Y}_\theta^s)\|\mathsf{Y}_\theta^s\|_2\right)\right] \qquad (10)
\end{aligned}
$$

For a pair $(l,\ s)$ of positive integers, we define the following event, over the probability space defined by G:

$$
\begin{aligned}
&\Omega(l,\ s)\\
&:= \left\{\left\|\mathsf{X}_\theta^l\left(\mathsf{G}^\top\mathsf{G}-\mathsf{I}\right)\mathsf{Y}_\theta^s\right\|_2 \geq \left(\varepsilon + c_2\frac{r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)}{c_1\varepsilon t}\right)\left\|\mathsf{X}_\theta^l\right\|_2\|\mathsf{Y}_\theta^s\|_2\right\},
\end{aligned}
$$

where $c_1,\ c_2 > 0$ are positive constants. According to Lemma 1, the above event holds with probability:

$$
\mathbb{P}\left(\Omega(l,s)\right) \leq c_2^{r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)}\cdot e^{-c_1\gamma_{l,s}t\cdot\min(\gamma_{l,s},1)} \qquad (8)
$$

where $\gamma_{l,s} = \varepsilon + c_2\frac{r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)}{c_1\varepsilon t}$. A useful observation for (8) is given in the next lemma; the proof is provided in the appendix.

Lemma 2. *Fix integer $t > 0$. The following inequality holds*

$$
\mathbb{P}\left(\Omega(l,s)\right) \leq e^{-c_1 t\cdot\min(\varepsilon,\varepsilon^2)}, \quad \forall l, s. \qquad (9)
$$

By conditioning on the event $\bigcap_{s=1}^{S}\bigcap_{l=1}^{L}\Omega(l,s)^c$, we can upper bound (7) as in (10) where the infinite series $\sum_{i=1}^{\infty}e^{-i+1} = \frac{e}{e-1}$ is used in the last equality.

To proceed, we observe the following for $\sum_{l=1}^{L}r(\mathsf{X}_\theta^l)\left\|\mathsf{X}_\theta^l\right\|_2$; similar reasoning applies for $\sum_{s=1}^{S}r(\mathsf{Y}_\theta^s)\|\mathsf{Y}_\theta^s\|_2$. By construction in (iii):

$$
\left\|\mathsf{X}_\theta^l\right\|_2 < e^{-l+1} = e\cdot e^{-l} \leq e\cdot\min_{i\in\mathcal{I}_l}\sigma_i.
$$

Thus, it is obvious that:

$$
r(\mathsf{X}_\theta^l)\cdot\left\|\mathsf{X}_\theta^l\right\|_2 < e\cdot r(\mathsf{X}_\theta^l)\cdot\min_{i\in\mathcal{I}_l}\sigma_i \leq e\cdot\|\mathsf{X}_\theta^l\|_\star,
$$

since $\|\mathsf{E}\|_\star \geq r(\mathsf{E})\cdot\min_j\sigma_j(\mathsf{E})$, for any matrix E. To this end,

$$
\sum_{l=1}^{L}r(\mathsf{X}_\theta^l)\left\|\mathsf{X}_\theta^l\right\|_2 < e\sum_{l=1}^{L}\|\mathsf{X}_\theta^l\|_\star = e\|\mathsf{X}_\theta\|_\star.
$$

Similarly, we have $\sum_{s=1}^{S}r(\mathsf{Y}_\theta^s)\|\mathsf{Y}_\theta^s\|_2 < e\|\mathsf{Y}_\theta\|_\star$ and, therefore, (10) becomes:

$$
\begin{aligned}
\left\|\mathsf{X}_\theta\left(\mathsf{G}^\top\mathsf{G}-\mathsf{I}\right)\mathsf{Y}_\theta\right\|_2 &< \frac{\varepsilon\cdot e^2}{(e-1)^2} + \frac{c_2}{c_1 t}\cdot\frac{e}{e-1}\left(\|\mathsf{X}\|_\star + \|\mathsf{Y}\|_\star\right)\\
&= \frac{\varepsilon\cdot e^2}{(e-1)^2} + \frac{c_2}{c_1 t}\cdot\frac{e}{e-1}\left(\mathbf{nr}\left(\mathsf{X}\right) + \mathbf{nr}\left(\mathsf{Y}\right)\right)\\
&\leq 3\varepsilon \qquad (11)
\end{aligned}
$$

where the last equality is satisfied since $t \geq \theta$ ($c_3 \geq \frac{c_2 e}{c_1(e-1)}$). By the union bound and since $0 < \varepsilon < 1$, (11) is violated with probability:

$$
\begin{aligned}
\mathbb{P}\left(\left(\bigcap_{s=1}^{S}\bigcap_{l=1}^{L}\Omega(l,s)^c\right)^c\right) &\overset{(9)}{\leq} \sum_{s=1}^{S}\sum_{l=1}^{L}e^{-c_1\cdot\varepsilon^2\cdot t}\\
&= e^{-c_1\cdot\varepsilon^2\cdot t + \ln\left(\left\lceil\ln\left(\frac{e}{\varepsilon}\right)\right\rceil\right)} \leq \frac{\delta}{2}
\end{aligned}
$$

where the last inequality is satisfied since $t \geq \frac{\ln(2/\delta)+\ln(\lceil \ln(e/\varepsilon) \rceil)}{c_1 \cdot \varepsilon^2}$.

The terms in (4) can be bounded as follows:

$$\begin{aligned}
&\|\mathsf{X}_\theta^c \mathsf{Y}_\theta\|_2 + \|\mathsf{X}_\theta \mathsf{Y}_\theta^c\|_2 + \|\mathsf{X}_\theta^c \mathsf{Y}_\theta^c\|_2 \\
&\qquad \leq \|\mathsf{X}_\theta^c\|_2 \|\mathsf{Y}_\theta\|_2 + \|\mathsf{X}_\theta\|_2 \|\mathsf{Y}_\theta^c\|_2 + \|\mathsf{X}_\theta^c\|_2 \|\mathsf{Y}_\theta^c\|_2 \\
&\qquad \leq 2\varepsilon + \varepsilon^2 \leq 3\varepsilon
\end{aligned}$$

where we used the fact that $\|\mathsf{X}_\theta\|_2 ,\ \|\mathsf{Y}_\theta\|_2 \leq 1$ and the inequalities $\|\mathsf{X}_\theta^c\|_2 \leq \varepsilon$ (similarly for Y_θ^c).

To bound the terms in (3), we provide the next Lemma; the proof is given in the appendix.

Lemma 3. *If $t \geq \theta + 8\ln(8/\delta)$ and $0 < \varepsilon < 1$, then, with probability at least $\delta/2$, the term appearing in* (3) *is at most* 33ε.

Applying the union bound on Lemma 3 and the complement of $\bigcap_{s=1}^{S} \bigcap_{l=1}^{L} \Omega(l,s)^c$, we conclude that the following inequality holds:

$$\left\| \mathsf{X}\mathsf{G}^\top \mathsf{G}\mathsf{Y} - \mathsf{X}\mathsf{Y} \right\|_2 \leq 39\varepsilon,$$

with probability at least $1-\delta$. By rescaling ε we obtain the required result.

5 Application to data-driven low-dimensional embedding

As an application of the result above, we consider the following question:

PROBLEM: *Given a collection n points in $\mathbb{R}^d$, forming a matrix $\mathsf{A} \in \mathbb{R}^{n\times d}$ and an error parameter $\varepsilon > 0$, construct efficient and approximately accurate low-dimensional embedding $\mathsf{G} \in \mathbb{R}^{t\times d}$ such that:*

$$\|\mathsf{G}(\mathbf{a}_i - \mathbf{a}_j)^T\|_2^2 \overset{\pm\varepsilon}{\approx} \|\mathbf{a}_i - \mathbf{a}_j\|_2^2, \quad \forall \mathbf{a}_i \in \mathbb{R}^{1\times d},\ \mathbf{a}_i \in \mathsf{Rows}(\mathsf{A}).$$

The target dimension t is to heavily depend on the input data matrix A and be independent of the input dimensions.

Inspired by Theorem 1, the following theorem proposes a data-dependent randomized low-dimensional embedding G with the following guarantees:

Theorem 2. *Let $\mathsf{A} \in \mathbb{R}^{n\times d}$. Moreover, assume that A with nuclear rank $\boldsymbol{nr}(\mathsf{A}) \ll O(\log n)$. If $t = \Omega\left(\frac{\boldsymbol{nr}(\mathsf{A})+\log(\log(1/\varepsilon))+\log(1/\delta)}{\varepsilon^2}\right)$, then the following inequalities hold*

$$\forall i,j \in [n] \quad \left| \left\| \mathsf{G}\left(\mathbf{a}_i - \mathbf{a}_j\right)^\top \right\|_2^2 - \|\mathbf{a}_i - \mathbf{a}_j\|_2^2 \right| \leq 2\varepsilon \|\mathsf{A}\|_2^2,$$

with probability at least $1-\delta$.

To fully specify the algorithmic procedure followed in practice, only $\mathsf{A} \in \mathbb{R}^{n\times d}$ and constants $\delta,\ \varepsilon \in (0,1)$ are given as input and $\mathsf{G} \in \mathbb{R}^{t\times d}$ is returned. Moreover, G guarantees to preserve the distances between the rows of A with additive error $2\varepsilon\|\mathsf{A}\|_2^2$ and with probability at least $1-\delta$. The exact steps followed are given in Algorithm 1.

Algorithm 1 Data-Driven Random Linear Embedding

Input: $\mathsf{A} \in \mathbb{R}^{n\times d}, \varepsilon, \delta \in (0,1)$

1: Compute (or approximate) $\|\mathsf{A}\|_\star$ and $\|\mathsf{A}\|_2$.
2: Set parameter $t = \Omega\left(\frac{\mathbf{nr}(\mathsf{A})+\log\log(1/\varepsilon)+\log(1/\delta)}{\varepsilon^2}\right)$.
3: Generate $\mathsf{G} \in \mathbb{R}^{t\times n}$ according to Theorem 2.
4: **Output:** Matrix G and GA that satisfies Theorem 2 with probability at least $1-\delta$.

Proof. By substituting $\mathsf{X} = \mathsf{Y}^\top = \mathsf{A}$ in Theorem 1, we have:

$$\mathbb{P}\left(\left\| \mathsf{A}\mathsf{G}^\top\mathsf{G}\mathsf{A}^\top - \mathsf{A}\mathsf{A}^\top \right\|_2 \leq \varepsilon \|\mathsf{A}\|_2^2 \right) \geq 1-\delta.$$

By the definition of the spectral norm, the condition above can be further written as the following maximization problem:

$$\max_{\mathbf{x}:\mathbf{x}\in\mathbb{R}^n, \|\mathbf{x}\|_2=1} \left\{ \left| \mathbf{x}^\top \left(\mathsf{A}\mathsf{G}^\top\mathsf{G}\mathsf{A}^\top - \mathsf{A}\mathsf{A}^\top \right) \mathbf{x} \right| \right\} \leq \varepsilon \|\mathsf{A}\|_2^2 .$$

Moreover, it is obvious that, if we restrict the search space $\mathbf{x} \in \mathbb{R}^n$ to $\mathbf{x} \in \mathbb{B}^n$ where:

$$\mathbb{B}^n = \left\{ \mathbf{y} \mid \mathbf{y} = \frac{1}{\sqrt{2}} \left(\mathbf{e}_i - \mathbf{e}_j \right), \forall i, j \in [n] \right\},$$

where $\mathbf{e}_i$ is the standard basis vector with 1 in the i-th position, we further have:

$$\begin{aligned} &\max_{\mathbf{x}:\mathbf{x}\in\mathbb{B}^n} \left\{ \left| \mathbf{x}^\top \left(\mathsf{A}\mathsf{G}^\top \mathsf{G}\mathsf{A}^\top - \mathsf{A}\mathsf{A}^\top \right) \mathbf{x} \right| \right\} \\ &\qquad \leq \max_{\mathbf{x}:\mathbf{x}\in\mathbb{R}^n, \|\mathbf{x}\|_2 = 1} \left\{ \left| \mathbf{x}^\top \left(\mathsf{A}\mathsf{G}^\top \mathsf{G}\mathsf{A}^\top - \mathsf{A}\mathsf{A}^\top \right) \mathbf{x} \right| \right\} \end{aligned}$$

which leads to:

$$\mathbb{P}\left(\max_{\mathbf{x}:\mathbf{x}\in\mathbb{B}^n} \left\{ \left| \mathbf{x}^\top \left(\mathsf{A}\mathsf{G}^\top \mathsf{G}\mathsf{A}^\top - \mathsf{A}\mathsf{A}^\top \right) \mathbf{x} \right| \right\} \leq \varepsilon \left\| \mathsf{A} \right\|_2^2 \right) \geq 1 - \delta. \tag{12}$$

Observe also that $\left| \mathbf{x}^\top \left(\mathsf{A}\mathsf{G}^\top \mathsf{G}\mathsf{A}^\top - \mathsf{A}\mathsf{A}^\top \right) \mathbf{x} \right| = \left| \|\mathsf{G}\mathsf{A}^\top \mathbf{x}\|_2^2 - \|\mathsf{A}^\top \mathbf{x}\|_2^2 \right|$, which further transforms (12) as:

$$\mathbb{P}\left(\max_{\mathbf{x}:\mathbf{x}\in\mathbb{B}^n} \left\{ \left| \|\mathsf{G}\mathsf{A}^\top \mathbf{x}\|_2^2 - \|\mathsf{A}^\top \mathbf{x}\|_2^2 \right| \right\} \leq \varepsilon \|\mathsf{A}\|_2^2 \right) \geq 1 - \delta. \tag{13}$$

For any vector $\mathbf{x} \in \mathbb{B}^n$, we observe that $\mathbf{x}^\top \mathsf{A}\mathsf{G}^\top = 1/\sqrt{2} \left(\mathbf{a}_i - \mathbf{a}_j \right) \mathsf{G}^\top$ while $\mathsf{G}\mathsf{A}^\top \mathbf{x} = 1/\sqrt{2} \mathsf{G} \left(\mathbf{a}_i - \mathbf{a}_j \right)^\top$. Similarly, $\mathbf{x}^\top \mathsf{A} = 1/\sqrt{2} \left(\mathbf{a}_i - \mathbf{a}_j \right)$ and $\mathsf{A}^\top \mathbf{x} = 1/\sqrt{2} \left(\mathbf{a}_i - \mathbf{a}_j \right)^\top$. Thus, (13) becomes:

$$\begin{aligned} \mathbb{P}\Big(\max_{i,j\in[d], i\neq j} \left\{ \left| \left\| \mathsf{G} \left(\mathbf{a}_i - \mathbf{a}_j \right)^\top \right\|_2^2 - \|\mathbf{a}_i - \mathbf{a}_j\|_2^2 \right| \right\} \leq 2\varepsilon \|\mathsf{A}\|_2^2 \Big) \\ \geq 1 - \delta. \end{aligned} \tag{14}$$

Since (14) is satisfied for the maximizing combination of i, j, we can safely remove the maximization to get:

$$\mathbb{P}\left(\left| \left\| \mathsf{G} \left(\mathbf{a}_i - \mathbf{a}_j \right)^\top \right\|_2^2 - \|\mathbf{a}_i - \mathbf{a}_j\|_2^2 \right| \leq 2\varepsilon \|\mathsf{A}\|_2^2 \right) \geq 1 - \delta,$$

which completes the proof. □

A complete set of experiments will be included in an extended version of the paper.

6 Conclusions

We present a novel analysis for a class of randomized and provably ε-accurate algorithms for the problem of matrix multiplication. As an application of this result, we show the utilization of the proposed scheme on data-driven low dimensional embeddings with additive error approximation, in the case where the data live on a subspace characterized by a small *nuclear rank*.

An interesting question to pursue lies in the substitution of nuclear rank by stable rank: recent developments on this topic [8] show similar results using the latter metric as the intrinsic data dimension; a weaker assumption than the nuclear rank. However, the dependence on the approximation error is of the order $O(\frac{1}{\varepsilon^4})$, as opposed to $O(\frac{1}{\varepsilon^2})$ presented in this work. We hope this paper triggers future efforts to improve stable rank-based bounds with respect to error dependency.

Acknowledgment

A. Zouzias would like to thank Mark Rudelson for several discussions on the approximate matrix multiplication problem. This research has received funding from the ERC under the EU's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement n^o 259569.

References

[1] T. Sarlos, "Improved approximation algorithms for large matrices via random projections," in *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*. IEEE, 2006, pp. 143–152.

[2] V. Chandrasekaran and M. I. Jordan, "Computational and statistical tradeoffs via convex relaxation," *Proceedings of the National Academy of Sciences*, vol. 110, no. 13, pp. E1181–E1190, 2013.

[3] D. Achlioptas, Z. Karnin, and E. Liberty, "Near-optimal entrywise sampling for data matrices," *arXiv preprint arXiv:1311.4643*, 2013.

[4] M. W. Mahoney and P. Drineas, "Cur matrix decompositions for improved data analysis," *Proceedings of the National Academy of Sciences*, vol. 106, no. 3, pp. 697–702, 2009.

[5] C. Boutsidis, M. W. Mahoney, and P. Drineas, "An improved approximation algorithm for the column subset selection problem," in *SODA*, 2009, pp. 968–977.

[6] P. Indyk and R. Motwani, "Approximate nearest neighbors: towards removing the curse of dimensionality," in *Proceedings of the thirtieth annual ACM symposium on Theory of computing*. ACM, 1998, pp. 604–613.

[7] C. Boutsidis, A. Zouzias, and P. Drineas, "Random projections for k-means clustering," in *NIPS*, 2010, pp. 298–306.

[8] A. Zouzias, "Randomized primitives for linear algebra and applications," Ph.D. dissertation, University of Toronto, 2013.

[9] A. Magen and A. Zouzias, "Low rank matrix-valued chernoff bounds and approximate matrix multiplication," in *SODA*, 2011, pp. 1422–1436.

[10] W. B. Johnson and J. Lindenstrauss, "Extensions of lipschitz mappings into a hilbert space," *Contemporary mathematics*, vol. 26, no. 189-206, p. 1, 1984.

[11] R. G. Baraniuk and M. B. Wakin, "Random projections of smooth manifolds," *Foundations of Computational Mathematics*, vol. 9, no. 1, pp. 51–77, 2009. [Online]. Available: http://dx.doi.org/10.1007/s10208-007-9011-z

[12] K. L. Clarkson, "Tighter bounds for random projections of manifolds," in *Proceedings of the twenty-fourth annual symposium on Computational geometry*. ACM, 2008, pp. 39–48.

[13] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann, "Reconstruction and subgaussian operators in asymptotic geometric analysis," *Geometric and Functional Analysis*, vol. 17, no. 4, pp. 1248–1282, 2007. [Online]. Available: http://dx.doi.org/10.1007/s00039-007-0618-7

[14] B. Klartag and S. Mendelson, "Empirical processes and random projections," *Journal of Functional Analysis*, vol. 225, no. 1, pp. 229 – 245, 2005. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0022123604003635

[15] P. K. Agarwal, S. Har-Peled, and H. Yu, "Embeddings of surfaces, curves, and moving points in euclidean space," in *Proceedings of the Twenty-third Annual Symposium on Computational Geometry*, ser. SCG '07. New York, NY, USA: ACM, 2007, pp. 381–389. [Online]. Available: http://doi.acm.org/10.1145/1247069.1247135

[16] Y. Bartal, B. Recht, and L. J. Schulman, "Dimensionality reduction: Beyond the johnson-lindenstrauss bound," in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '11. SIAM, 2011, pp. 868–887. [Online]. Available: http://dl.acm.org/citation.cfm?id=2133036.2133104

[17] L.-A. Gottlieb and R. Krauthgamer, "A nonlinear approach to dimension reduction," in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '11. SIAM, 2011, pp. 888–899. [Online]. Available: http://dl.acm.org/citation.cfm?id=2133036.2133105

[18] P. Indyk and A. Naor, "Nearest-neighbor-preserving embeddings," *ACM Trans. Algorithms*, vol. 3, no. 3, Aug. 2007. [Online]. Available: http://doi.acm.org/10.1145/1273340.1273347

[19] C. Hegde, A. Sankaranarayanan, W. Yin, and R. Baraniuk, "A convex approach for learning near-isometric linear embeddings," *preparation, August*, 2012.

[20] E. Grant, C. Hegde, and P. Indyk, "Nearly optimal linear embeddings into very low dimensions," *IEEE GlobalSIP Symposium on Sensing and Statistical Inference, Austin, TX*, 2013.

[21] M. Rudelson and R. Vershynin, "Hanson-wright inequality and sub-gaussian concentration," *Electron. Commun. Probab.*, vol. 18, pp. no. 82, 1–9, 2013.

Appendix

In our analysis, we make use of the following concentration bound on the operator norm of the product of a fixed matrix with a Gaussian matrix which is a direct consequence of concentration of Lipschitz function on Gaussian space, see e.g. [8, p. 10].

Lemma 4. *Let* $\mathsf{X} \in \mathbb{R}^{n\times d}$ *and* $t \geq 1$. *For every* $\tau > 0$*:*

$$\mathbb{P}\left(\left\|\mathsf{X}\mathsf{G}^\top\right\|_2 \geq \|\mathsf{X}\|_{\mathrm{F}}/\sqrt{t} + \|\mathsf{X}\|_2 + \tau\|\mathsf{X}\|_2/\sqrt{t}\right) \leq \exp(-\tau^2/8).$$

6.1 Proof of Lemma 2

In all the cases below, observe $\gamma_{l,s} > \varepsilon$ by definition. Using Eqn. (9), we consider the following two cases:

(i) $\gamma_{l,s} < 1$: in this case, we have $\min(\gamma_{l,s}, 1) \geq \min(\varepsilon, 1) = \varepsilon$. Thus, (9) becomes:

$$\begin{aligned}\mathbb{P}\left(\Omega(l,s)\right) &\leq c_2^{r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)} \cdot e^{-c_1\left(\varepsilon + c_2\frac{r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)}{c_1\varepsilon t}\right)t\cdot\varepsilon} \\ &\leq e^{-c_1\varepsilon^2 t} \cdot e^{-c_2\left(r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)\right)} \cdot e^{\ln(c_2)\left(r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)\right)} \\ &\leq e^{-c_1\varepsilon^2 t} \end{aligned} \tag{15}$$

(ii) $\gamma_{l,s} \geq 1$: in this case, $\min(\gamma_{l,s}, 1) = 1$. Following the same steps as above, we have:

$$\begin{aligned}\mathbb{P}\left(\Omega(l,s)\right) &\leq c_2^{r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)} \cdot e^{-c_1\left(\varepsilon + c_2\frac{r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)}{c_1\varepsilon t}\right)t} \\ &\leq e^{-c_1\varepsilon t} \cdot e^{\left(1-\frac{1}{\varepsilon}\right)c_2\left(r(\mathsf{X}_\theta^l)+r(\mathsf{Y}_\theta^s)\right)}.\end{aligned}$$

However, since $0 < \varepsilon < 1$, we have $1 - \frac{1}{\varepsilon} < 0$ and thus, the above inequality can be further upper bounded by:

$$\mathbb{P}\left(\Omega(l,s)\right) \leq e^{-c_1\varepsilon t}.$$

6.2 Proof of Lemma 3

Apply Lemma 4 with $\tau = \sqrt{t}$ on all four matrices X_θ^c, X_θ, Y_θ and Y_θ^c. With probability at least $1 - \exp(-t/8)$:

$$\begin{aligned}\left\|\mathsf{X}_\theta^c\mathsf{G}^\top\right\|_2 &\leq \frac{\|\mathsf{X}_\theta^c\|_{\mathrm{F}}}{\sqrt{t}} + \left(1+\frac{\tau}{\sqrt{t}}\right)\|\mathsf{X}_\theta^c\|_2 \\ &\overset{\tau=\sqrt{t}}{=} \frac{1}{\sqrt{t}}\|\mathsf{X}_\theta^c\|_{\mathrm{F}} + 2\varepsilon \\ &\leq \frac{1}{\sqrt{t}}\|\mathsf{X}\|_{\mathrm{F}} + 2\varepsilon \leq \frac{1}{\sqrt{t}}\cdot\|\mathsf{X}\|_\star + 2\varepsilon \\ &\overset{t\geq\theta}{\leq} \frac{1}{\sqrt{\theta}}\mathbf{nr}\left(\mathsf{X}\right) + 2\varepsilon \leq \frac{\varepsilon}{\sqrt{c_3}} + 2\varepsilon \overset{c_3>1}{\leq} 3\varepsilon\end{aligned}$$

by definition of θ and $\|\mathsf{X}_\theta^c\|_2 \leq \varepsilon$. Similarly, with probability at least $1 - \exp(-t/8)$:

$$\begin{aligned}\left\|\mathsf{X}_\theta\mathsf{G}^\top\right\|_2 &\leq \frac{\|\mathsf{X}_\theta\|_{\mathrm{F}}}{\sqrt{t}} + \left(1+\frac{\tau}{\sqrt{t}}\right)\|\mathsf{X}_\theta\|_2 \\ &\overset{\tau=\sqrt{t}}{=} \frac{1}{\sqrt{t}}\|\mathsf{X}_\theta\|_{\mathrm{F}} + 2 \\ &\leq \frac{1}{\sqrt{t}}\|\mathsf{X}\|_{\mathrm{F}} + 2 \leq \frac{1}{\sqrt{t}}\cdot\|\mathsf{X}\|_\star + 2 \\ &\overset{t\geq\theta}{\leq} \frac{1}{\sqrt{\theta}}\mathbf{nr}\left(\mathsf{X}\right) + 2 \leq \frac{\varepsilon}{\sqrt{c_3}} + 2 \overset{c_3>1}{\leq} 2(\varepsilon+1)\end{aligned}$$

by definition of θ and $\|\mathsf{X}_\theta\|_2 = 1$. We further observe $\left\|\mathsf{Y}_\theta^c\mathsf{G}^\top\right\|_2 \leq 3\varepsilon$ and $\left\|\mathsf{X}_\theta\mathsf{G}^\top\right\|_2 \leq 2(\varepsilon+1)$ both with probability at least $1 - \exp(-t/8)$. Union bound all the above four results, it follows that

$$\begin{aligned}&\left\|\mathsf{X}_\theta^c\mathsf{G}^\top\mathsf{G}\mathsf{Y}_\theta\right\|_2 + \left\|\mathsf{X}_\theta\mathsf{G}^\top\mathsf{G}\mathsf{Y}_\theta^c\right\|_2 + \left\|\mathsf{X}_\theta^c\mathsf{G}^\top\mathsf{G}\mathsf{Y}_\theta^c\right\|_2 \\ &\leq 12\varepsilon\left(\varepsilon+1\right) + 9\varepsilon^2 \leq 12\varepsilon + 21\varepsilon^2 \leq 33\varepsilon\end{aligned}$$

which holds with probability at most $4\cdot e^{-t/8}$. By definition of t, the above inequality is violated with probability at most $\delta/2$.